LAWRENCE LIVERMORE NATIONAL LABORATORY

UCRL-TR-234773

CASC progress Notes August 2007

Xabier Garaizar

September 19, 2007

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

August 2007

The Center for Applied Scientific Computing (CASC) applies computational expertise to problems of national importance at Lawrence Livermore National Laboratory (LLNL). CASC staff members collaborate actively with other discipline experts at LLNL, other national laboratories, academia, and industry. The purpose of these notes is to provide information about some of the current research and advanced development activities in CASC, with special emphasis on DNT relevance. Interested readers having any questions should contact the appropriate contributor or the CASC director, John Grosh (4-6520, grosh1@llnl.gov). Contact information is given for the main CASC point of contact for each of the projects listed.

Scalable Linear Solvers

Rob Falgout
rfalgout@llnl.gov
2-4377

The goal of the Scalable Linear Solvers (SLS) project is to develop scalable algorithms and software for solving large, sparse linear systems of equations on parallel computers. The primary software product is *hypre*, a library of high performance preconditioners and solvers that features parallel multigrid methods for both structured and unstructured grid problems. The *hypre* software is used by major codes in both AX and B Divisions.

Selected highlights

- ASC L2 Milestone. The SLS project recently completed an ASC L2 milestone entitled "Performance assessment and improvements of *hypre* in ASC codes". The purpose of this milestone was to complete performance studies and demonstrate improved performance of the *hypre* linear solver library. The simulations of interest for these studies were chosen in conjunction with several code teams across the A-, B-, and ICF-Programs. The vast majority of code and algorithmic improvements derived from these simulations are now available in public versions of the *hypre* library. Through the use of new algorithms, improvements to existing algorithms, and improvements to parameter settings for existing algorithms, the *hypre* team demonstrated significant performance improvements (varying from 1.2× to 25× faster) by measuring time to solution in several ASC simulation codes.
- AMS Solver. The Auxiliary-space Maxwell Solver (AMS) was recently developed and added to the *hypre* linear solver library. AMS is designed to solve the definite Maxwell's equations which arise in many electromagnetic or magneto-hydrodynamic (MHD) simulations. This new method utilizes projections onto a series of carefully-designed spaces to transform the difficult problem of applying multigrid techniques directly to the definite Maxwell equations to that of solving a sequence of well-understood problems for which multigrid methods already exist. Because of this, AMS is able to build on the BoomerAMG algebraic multigrid solver in *hypre*. The new solver is the first provably scalable solver for electromagnetic problems on general unstructured meshes and is being actively used as an MHD solver by the ASC code ALE3D.
- Long-range interpolation**.** The algebraic multigrid algorithm BoomerAMG in *hypre* is being used by most of the ASC codes at LLNL. The parallel scalability of this algorithm is crucial, but difficult because algorithmic complexity is unpredictable and difficult to control. New parallel coarse-grid selection algorithms were recently developed to improve complexity and greatly reduce solution time at large processor scales, but solver convergence properties suffered. To recover ideal convergence rates while retaining the improved complexity benefits, new long-range interpolation techniques were

developed. These interpolation techniques have shown as much as a factor of 4.5× speedup in solution time in addition to the speedups gained by the new parallel coarsening algorithms. This improved algorithm is now benefiting the ASC codes and contributed greatly to the successful L2 Milestone mentioned above.

- <u>R&D100 Award</u>**.** The *hypre* team won an R&D100 Award, one of five won by the Lab this year. These awards are considered the "Oscars of Invention" and only 100 are awarded nationally each year. The *hypre* software library is unique in its ability to provide solution algorithms that are effective on a spectrum of diverse problems, are easily accessible via multiple user interfaces, and effectively utilize the computer power of today's supercomputers to perform larger simulations.

Phase Field Modeling

Milo Dorr, Jim Belak (PAT), Patrice Turchi (CMLS)
milodorr@llnl.gov
4-6013

The goal of this LDRD-supported project is to develop a new capability for simulating phase and microstructure evolution in polycrystalline materials. The relevance to DNT is the prospect of improved materials models in ASC codes, enabling more accurate representations of dynamic phase transitions in metals and alloys under extreme conditions. Although molecular dynamics (MD) simulations are already capable of providing detailed atomistic predictions of such processes, they are computationally limited to scales of a few hundred nanometers and a couple tenths of a nanosecond. Through the use of a continuous description of phase parameters evolved by a system of equations solved with high-performance algorithms, the promise of phase field modeling is to extend the MD-initiated microstructure predictions by a few orders of magnitude in space and time, *i.e.*, the hydrodynamic scales of interest to DNT.

<u>*Progress and current status*</u>:
We have developed a research code called AMPE (Adaptive Mesh Phase Evolution) to solve a 3D system of Ginzburg-Landau equations for a phase parameter and grain orientations described by quaternions. The use of adaptive mesh refinement enables the computational mesh to be concentrated at the grain boundaries for efficiency. AMPE is fully parallel and is based on SAMRAI (Structured Adaptive Mesh Refinement Infrastructure). Since the integration of the grain orientation quaternions is necessarily implicit, we also employ linear solvers from Hypre and Sundials. In addition, we have developed a methodology for converting MD data computed by the Streitz/Glosli DDCMD code to phase field data, which is then used to provide an initial condition for AMPE. We are currently beginning the process of validating AMPE against MD simulations.

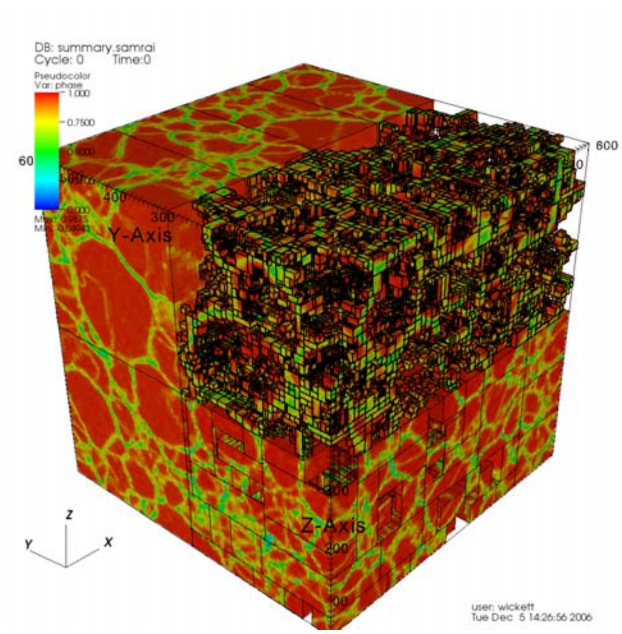

The NIF-ALE-AMR Project and Fragmentation Modeling

Alice Koniges
koniges@llnl.gov
3-7890

Fragmentation is a fundamental material process that naturally spans spatial scales from microscopic to macroscopic, with applications in a variety of LLNL programs. We are developing new models for fragmentation that can be run using an innovative combination of hierarchical material modeling (HMM), adaptive mesh refinement (AMR) and arbitrary Lagrangian Eulerian (ALE) methods. We have developed a new standalone code aimed at making predictions of the debris and shrapnel generated by NIF experiments using these new technologies. We have also conducted a series of dedicated experiments on thin foil fragmentation and collected data for validation of our new models.

Recent project highlights:

We have developed a new fragmentation model that uses void insertion to follow the development of cracks and the formation of fragments and applied this model to several problems of interest including fragment size prediction in expanding ring geometries and thin foil ablation and spall.

We have connected our simulations to the LLNL MS material libraries developed by Rich Becker and Nathan Barton. We have also implemented our own failure models for both brittle and ductile materials.

We have used the Voronoi method to generate a polycrystal implementation at one AMR level and an isotropic model at the higher AMR levels and are studying the implications of and methods for combining these different methods in the same simulation.

We implemented a new multimaterial interface reconstruction scheme on the 2D and 3D ALE-AMR meshes and are in the process of connecting our code with a radiation diffusion solver library we designed specifically for the ALE-AMR mesh.

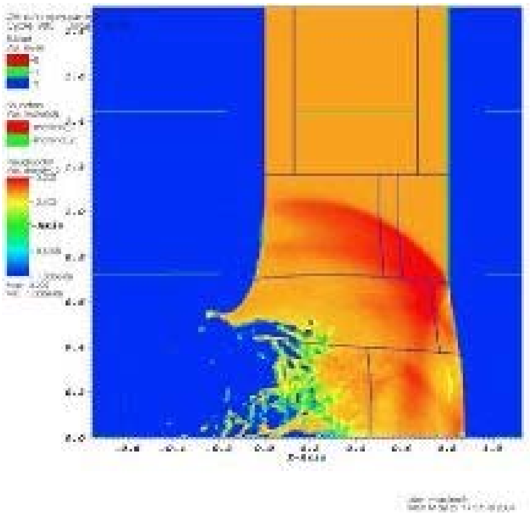

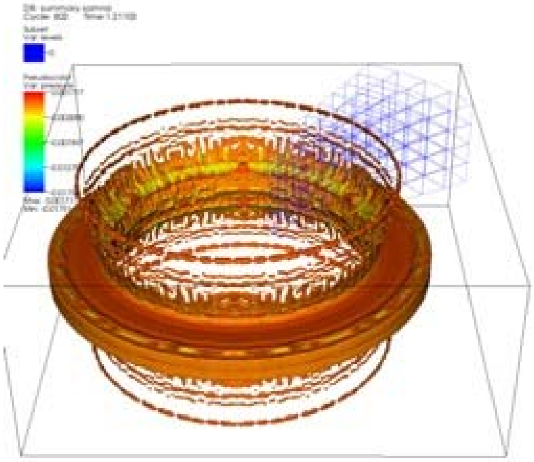

Projective Integrators for Multiphysics Simulations (PRISMS)

Steven Lee
lee117@llnl.gov
4-5989

Multiphysics simulations need to account for various, important physical effects across a wide range of space- and time-scales. In some cases, these subscale effects may come in the form of output from black-box codes -- which then needs to be seamlessly incorporated into the overall simulation. These black-boxes are often trusted components that stem from many years of development, testing and use.

For KULL, we have applied projective integration-based ideas for seamlessly coupling a material strength library (MSlib) with KULL's 3-stage Runge-Kutta time stepper (RK3). The RK3-MSlib integrator will enable KULL to incorporate effects from material hardening (or melting) in a manner that uses MSlib strictly as a black-box. Careful implementation and testing within KULL are underway. This research work is a collaboration among Steven Lee (CASC), Armand Attia and Paul Amala (AX-Division).

Recent highlights

We have formulated a novel technique for coupling RK3 and MSlib in a way that uses MSlib strictly in a black-box fashion. This approach makes 6 calls to MSlib, requiring one, two, and then three MSlib calls for the 3 successive substeps within a complete RK3 step.

We have proposed a second formulation which, given certain assumptions, requires only a total of 3 calls to MSlib for a complete RK3 step. These assumptions relate to strength behavior as a material hardens. The cost, accuracy, and range of validity, will vary for these formulations and verification studies are in progress.

The Mesquite Mesh Quality Improvement Toolkit

Lori Diachin
diachin2@llnl.gov
2-7130

Mesh quality is critical for accuracy and efficiency in the solution of PDE-based applications. Quality can be affected in many stages of the solution process from de-featuring CAD models, to mesh generation procedures and h- and r-adaptive schemes. There are a number of techniques that have been developed for mesh improvement ranging from simple Laplacian smoothing to more sophisticated algorithms such as Winslow smoothers for structured meshes and the numerical optimization methods and topology modification schemes recently developed for unstructured meshes.

As part of the ITAPS project in the DOE SciDAC program, we have been working with scientists at Sandia National Laboratory to develop the *Mesquite* mesh quality improvement toolkit. The primary aim of this project is to provide a freely available, comprehensive software package that accommodates a number of different mesh element types, quality metrics, high-level solution strategies, and mesh optimization algorithms. Mesquite does element shape improvement, mesh untangling, and deforming meshes on local mesh patches consisting of either triangular, tetrahedral, quadrilateral, hexahedral, or hybrid (including wedge and pyramid elements) unstructured meshes. It has a number of state-of-the-art algorithms for optimization-based node point movement including steepest descent, conjugate gradient, feasible Newton, and active set solvers. In addition, considerable attention was given to the development of a flat mesh data structure for the internal representation of unstructured mesh data that is both highly efficient and convenient.

Mesquite is designed to be easily extensible in order to support new research in mesh quality improvement and to allow rapid delivery of customized mesh optimization algorithms to applications. The use of C++ classes such as MeshQualityMetric, ObjectiveFunction, QualityAssessor, VertexMover, and InstructionQueue help us achieve an object-oriented, flexible design. Embedded within the various classes are member functions which perform the computationally intensive numerical calculations. These member functions avoid objects, use pointers, arrays, and other low-level data structures to ensure that the computations are as efficient as possible.

In the past year we have worked to incorporate Mesquite algorithms into ALE3d. Initial tests have shown that the optimization-based techniques in Mesquite are effective in improving mesh quality as the mesh moves to follow application physics. Ongoing work is focused on developing a parallel version of Mesquite and improving single processor performance. We are also working to combine the Mesquite smoothing algorithms with a stand-alone library that provides face and edge swapping for simplicial meshes.

KULL

Michael Lambert
mlambert@llnl.gov
3-6642

KULL is an AX Division parallel 2D and 3D unstructured mesh radiation hydrodynamics code for ICF analysis. Several CASC scientists, matrixed to the KULL project, work on a variety of production level computer science, computational physics, and applied math issues.

- Michael Lambert brought more capabilities to the KULL IMC photon transport package. IMC Compton scattering physics recently passed the testing phase. IMC random walk acceleration performance gains were documented and peer reviewed. An IMC Level 3 milestone was completed. Along the way, common zone-boundary tracking for particle-based packages was implemented; this both reduced code maintenance and delivered lossless photon tracking through slightly pathological meshes with boomerang zones.

- Dan Laney advanced VisIt capabilities for KULL. He oversaw development of multigroup radiation field visualization using parallel axis plots. He oversaw completion of a graphics hardware accelerated HADES radiography code library (soon to be released), tested on KULL data sets. This demonstrated performance advantages for image analysis, and laid the foundation for an LDRD proposal. He also began an in-situ data visualization and querying tool that can attach to a running instance of KULL. Dan also coauthored a Best Application Paper at IEEE Visualization 2006: "Understanding the Structure of the Turbulent Mixing Layer in Hydrodynamic Instabilities".

- Brian Miller is now facilitating Sequoia procurement, part of which involves an updated UMT 2000 benchmark integrating a standalone scalable mesh generator, a mesh conversion layer, a version of the Kull Sn photon transport package, and a convenient Python/pyMPI front end. He and Brian McCandless also wrapped up scaling studies of the BG/L port of a subset of KULL, including some third party libraries, code infrastructure and the diffusion package – acceptable scaling was observed, largely influenced by Hypre, and performance was tuned by altering the mesh mapping onto interprocessor topology. In another effort, Brian and Patrick Brantley doubled the speed of an expensive physics package.

- Doug Peters performed ongoing integrated physics code-to-code comparisons for the KULL team. These stressed sensitivity of material and hydrodynamics algorithms. He reported earlier work along these lines at NECDC 2006. He continued his role as liaison to code users.

AX support

Ping Wang
wang32@llnl.gov
3-2612

Ping Wang supported NNSA Science Campaign 4, working on high fidelity 3D simulations relevant to energy balance. Radiation hydrodynamic simulations were used to begin the validation of subgrid scale models and to begin the incorporation of effects occurring at higher spatial dimensions into simulations at lower spatial dimensions. This high-profile work elegantly blends high resolution and faster calculation.

ROSE

Dan Quinlan
dquinlan@llnl.gov
3-2668

The ROSE project provides a way of moving a static code base to a wide range of new and future parallel architectures. ROSE makes it possible to write highly optimized software for large-scale scientific software applications targeting the modern special purpose state-of-the-art parallel computer architectures increasingly common within DOE. Compiler-based tools built using ROSE can be used to automate the introduction of domain-specific (e.g. MPI) performance optimizations into application codes. ROSE supports the development of source-to-source translators for C, C++, and both modern and legacy Fortran scientific applications. Our focus is on the types of optimizations either not yet addressed by common vendor compilers or too domain-specific to be attractive for vendors to support in their compilers (e.g. MPI, unstructured data optimizations, global analysis, etc.).

Recent highlights

- Recent work has added new program analysis to ROSE as part of research to optimize and verify correctness of MPI applications.

- We have also supported Jeff Keasler's new research work in B-div using ROSE to add a vector extension to C so that applications could be written at a higher level of abstraction and automatically restructured to map to low level hardware details.

- Empirical tuning for whole applications remains a current research topic and this year we were able to automate most of the steps in automatically tuning the SMG2000 code using 400+ different optimizations and selecting the one most appropriate for a specific target architecture.

- The latest work has built a static analysis tool, called *Compass*, for the evaluation of easily defined rules that enforce source code properties for the ALE3D and Kull projects; rules have been taken from ASC code project style guides and source codes can be automatically evaluated for violations of over 80 different rules (including domain-specific MPI usage).

AWARDS

2007 R&D 100 Awards

***hypre*.**
Rob Falgout, Panayot Vassilevski, Barry Lee, Ulrike Yang, Tzanio Kolev, Allison Baker, Jeff painter, Van Henson (CASC)

Distribution List:

C. Murray	L-001
D. Crawford	L-559
B. Goodwin	L-160
C. Verdon	L-013
M. Dunning	L-170
S. Ashby	L-561
M. McCoy	L-554
R. Christensen	L-160
J. Rathkopf	L-030
T. Adams	L-095
L. Kissel	L-554
T. Quinn	L-554
S. Louis	L-554
D. Pilkington	L-095
M. Zika	L-023
F. Graziani	L-095
R. Sharpe	L-151
R. Werne	L-176
D. McCallen	L-151
M. Dehghani	L-113
R. Couch	L-035
R. Ferencz	L-125
D. Nikkel	L-129
J. Sefcik	L-160
N. Burkhard	L-638
T. McAbee	L-095
B. Pudliner	L-095
D. Miller	L-095
G. Zimmerman	L-030
R. Pember	L-039
D. Nielsen	L-555
K. Budil	L-170
L. Terminello	L-125
D. Griswold	L-099
V. Bulatov	L-367
R. Hills	L-181
B. Cauble	L-041
K. Bradley	L-183
H. Scott	L-018
E. Dube	L-159
P. Folta	L-448
T. Damkroger	L-365
J. Grosh	L-561
X. Garaizar	L-561
J. McGraw	L-419
C. Matarazzo	L-047
L. Lagin	L-488
L. Diachin	L-550
S. Kohn	L-561
L. Howell	L-560